A large scale study of SVM based methods for abstract screening in systematic reviews*

Tanay Kumar Saha
Computer and Information Science
Indiana University-Purdue University
Indianapolis, IN 46202
Email: tksaha@iupui.edu

Mourad Ouzzani, Hossam M. Hammady and Ahmed K. Elmagarmid
Data Analytics Group
Qatar Computing Research Institute
Hamad Bin Khalifa University
Email: {mouzzani, hhammady, aelmagarmid}@qf.org.qa

arXiv:1610.00192v2 [cs.IR] 13 Oct 2016

***Abstract*—A major task in systematic reviews is abstract screening, i.e., excluding, often hundreds or thousand of, irrelevant citations returned from a database search based on titles and abstracts. Thus, a systematic review platform that can automate the abstract screening process is of huge importance. Several methods have been proposed for this task. However, it is very hard to clearly understand the applicability of these methods in a systematic review platform because of the following challenges: (1) the use of non-overlapping metrics for the evaluation of the proposed methods, (2) usage of features that are very hard to collect, (3) using a small set of reviews for the evaluation, and (4) no solid statistical testing or equivalence grouping of the methods. In this paper, we use feature representation that can be extracted per citation. We evaluate SVM based methods (commonly used) on a large set of reviews (61) and metrics (11) to provide equivalence grouping of methods based on a solid statistical test. Our analysis also includes a strong variability of the metrics using 500x2 cross validation. While some methods shine for different metrics and for different datasets, there is no single method that dominates the pack. Furthermore, we observe that in some cases relevant (included) citations can be found after screening only 15-20% of them via a certainty based sampling. A few included citations present outlying characteristics and can only be found after a very large number of screening steps. Finally, we present an ensemble algorithm for producing a 5-star rating of citations based on their relevance. Such algorithm combines the best methods from our evaluation and through its 5-star rating outputs a more easy-to-consume prediction.**

I. Introduction

Randomized controlled trials (RCTs) are a key component of medical research and by far the best way of achieving results that can genuinely increase our knowledge about treatment effectiveness [1]. Because of the huge explosion of RCTs, it is very hard for individuals to glean evidence from all of them. Systematic reviews synthesize the results of more than one RCT, summarize the effects of their individual outcomes with a certain degree of confidence and provide answers about the effectiveness of a particular intervention [2]. Thus, systematic reviews provide current best evidence for policy and clinical decision-making.

Systematic review involves formulating a research question, searching in multiple biomedical databases, identifying relevant studies based on abstracts and titles (abstract screening) and then based on full texts of a subset thereof, assessing their methodological qualities, data extraction and synthesis, and finally reporting the conclusions on the review question [3]. However, the search precision is quite low and it usually returns a large number of citations, from hundreds to thousands. Systematic review authors have to go through the tedious and time-consuming task of screening these citations. Therefore, a systematic review platform that can automate the abstract screening process is fundamental in expediting the process of processing systematic reviews.

There are several challenges in building a systematic review platform for abstract screening including: (1) feature extraction should be fast, (2) the features should be readily available, (3) the learning and prediction algorithms should be very efficient, and (4) the algorithms should also be able to handle the problem of extreme data imbalance (due to the low search precision). In addition, the presentation of the citations should be based on their relevance, i.e. all the relevant citations should be ranked higher than the irrelevant ones. This is an instance of the bipartite ranking problem. Generating a clear bipartition is a quite difficult task [4] and a common measure of success in such scenario is the area under the ROC curve (AUC). However, the optimization techniques based on the AUC can be reduced to binary classification [5] methods under specific settings. In this paper, we evaluate several SVM based methods that either use a loss function such as $0-1$ (Classification), AUC, KLD, or quadratic mean error.

To address the first two constraints, quick feature extraction and availability, we focus in this paper on two classes of features, namely uni-bigram and Word2Vec features. Other features may be hard to obtain such as co-citations or are not always readily available such as MeSH terms. Uni-Bigram features do not have such issues. Another popular choice is LDA based topic features. However, LDA features need to be generated per review and with each addition of a new batch of citations, one needs to regenerate the topics. There are several practical concerns about how many topics should be generated and the efficiency of generating the LDA model. Therefore, we use Word2Vec based features as it captures the semantic similarity between words similarly to LDA and once the model is built (offline), the features are readily available. On the other hand, for the data-imbalance constraint, i.e., irrelevant

This work was done while the first author was doing internship at QCRI, HBKU, Qatar. All private reviews/datasets used in this paper was collected from the existing abstract screening platform Rayyan(http://rayyan.qcri.org/).

citations are more common than the relevant ones, we evaluate algorithms that have the potentials to address this problem, namely SVM^{cost} (cost-sensitive SVM) and SVM^{perf} (SVM with multivariate performance measure). This will allow us to get more insights on the applicability of these algorithmic approaches to tackle the data imbalance issue. We should also note that SVM-based approaches have less parameters to tune compared to others such as random forests.

It is worth mentioning that a recent systematic review of current approaches in abstract screening [6] reported that among the 69 publications it surveyed, only 22 report about $Recall$, 18 report about $Precision$, and 10 about AUC. The non-overlapping usage of these metrics makes it difficult to draw a conclusion about the different methods. For example, if a particular work only reports $Recall$ but not $Precision$, it is hard to understand its applicability. Moreover, recent works [7], [8] suggest studying the variability of the reported metrics and advocate using a large number of repetitions (500×2) to make the conclusions more reproducible. Table I characterizes most of the existing approaches in terms of their most important aspects. In particular, we observe that most of the studies have been performed on a small number of reviews, the evaluation is reported with different metrics, the cross validation uses a very small number of repetitions, and most of the approaches did not perform proper statistical significant tests.

In this paper, using a large set of systematic reviews, we conduct an extensive evaluation of several classification methods, along with different feature representations, which have the potential to address the above constraints. This would give us insights on the best approach to adopt for abstract screening in a practical systematic review platform. We also present such an approach as a consequence of our evaluation. More specifically, we make the following contributions:

1) We use a large sample of reviews (61, Most of them were collected from an existing abstract screening platform: Rayyan[1])
2) We evaluate 18 different methods and report on 11 different metrics
3) We perform a 500x2 cross validation
4) We apply a 2-factor ANOVA analysis with a paired t-test and group the equivalent methods.
5) We present an ensemble method that present prediction results through a 5-star rating method.

The remainder of this paper is organized as follows. In Section II, we discuss related work. In Section III, we detail the SVM based methods used in the evaluation. Section IV presents the experimental results. Section V details the proposed 5-star rating method. We conclude in Section VI.

II. Related Work

We divide the related work into four groups: 1) work on "Abstract Screening" methods, 2) discussion of "Data Imbalance" as it is one of the main issues in the abstract screening process, 3) "Active Learning" a a popular method to address data imbalance issue, and 4) "Linear Review", from the legal domain, which bears some similarity with abstract screening in systematic reviews.

[1] http://rayyan.qcri.org/

Aspect	Related Studies
Feature Space Representation	Uni, Bigram, MeSH [9], LDA [10], Uni + Cite [11]
Algorithm	SVM Based [12], [13], Others [14]
Cross Validation	5×2 [9], [11], [13], [15], [16], 10-fold [17]
No. of Reviews used	15 [11], [15], [16], 6 [12], 18 [9], 3 [18], 1 [17]
Statistical Testing	Post-hoc paired Wilcoxon test [9], Rank Group [11]
Reported Metric	AUC [9], [11], Utility [12], Burden [12], Yield [12], WSS [11], Precision [13], Recall [11], F1 [13]

TABLE I
A brief summary of recent contributions and their characteristics. For more details about algorithms used in systematic reviews please consult [6], [14], and for metric usage [6]

A. Abstract Screening

A large body of past research has focused on automating the abstract screening process [9], [11]–[13], [15], [16]. In terms of feature representation, most of the existing approaches [15], [16] use unigram, bigram, and MeSH (Medical Subject Headings). An alternative to the MeSH terms is extracting LDA based latent topics from the titles along with the abstracts and using them as features [12]. Other methods [11], [15] utilize external information as features, such as citations and co-citations. In terms of methods, SVM-based learning algorithms are commonly used [9], [13], [15]. According to a recent study [14], 13 different types of algorithms have been proposed in systematic reviews including SVM, Decision Trees, Naive Bayes, k-nearest neighbor, and neural networks. To the best of our knowledge, for abstract screening, past works did not use the structured output version of SVM (SVM^{perf}) with different loss functions. SVM^{perf} [19] can learn faster than SVM. For prediction, it only keeps feature weights and thus the prediction module becomes very fast. On the other hand, as Transductive learning [20] takes both the labeled and unlabeled citations into account, the corresponding learning algorithm is slower than SVM^{perf} [19].

B. Data imbalance

Data imbalance in supervised classification is a well studied problem [19]–[23]. Two different kinds of approaches have been proposed to solve it. The first focuses on designing loss functions: KLD [24], Quadratic Mean [25], Cost sensitive

classification [26], Mean Average Precision [27], Random Forest [28] with meta-cost, and AUC [19]. The second kind generates synthetic data for artificially balancing the ratio of labeled relevant and irrelevant citations. The example of such methods are Borderline-SMOTE [29], Safe-level-SMOTE [30], and oversampling of the minority class along with undersampling the majority class. The authors in [31]–[33] use probability calibration techniques. In this paper, we evaluate the algorithmic approaches rather than the data centric approaches or methods based on probability calibration since synthetic data generation is computationally intensive—generating data in extreme imbalance becomes quite difficult. Besides, for probability calibration we need a validation set that might not be always available in such kind of abstract screening process.

C. Active Learning

As systematic reviews are done in batches, the problem of abstract screening can be modeled as an instance of batch mode active learning. Online learning trains the classifier after adding every citation whereas Batch-mode active learning (BAL) does the same after adding batches of citations. However, BAL does not have any theoretical guarantee compared to online learning [34], [35]. The task of BAL is to select batches with informative samples (citations) that would help learn a better classification model. There are two popular methods to select samples: (1) Certainty based and (2) Uncertainty based. In certainty based methods, "most certain" samples are selected to train the classifier, while in uncertainty based method ,"most uncertain" ones are selected. In the uncertainty sampling-based methodologies, a large number of uncertainty metrics have been proposed; examples include entropy [36], smallest-margin [37], least confidence [38], committee disagreement [39], and version space reduction [34], [40]. Using "most uncertain" samples based on these metrics improve the quality of the classifier to find the best separating hyperplane, and thus to improve its accuracy in classifying new instances [12]. On the other hand, certainty based methods have also been shown to be effective to carry out active learning on imbalanced data sets, as demonstrated in [41].

D. Linear Review

A very similar review system that is popular among law firms is Technology assisted Linear Review (TAR) [42]–[46]. The main objectives of both review systems are very similar. TAR is used to save time of the attorneys to screen relevant documents rather than on the irrelevant ones. Generally the number of documents are huge (millions) compared to the same in systemtic review, and, so active learning becomes a popular method. However, unlike systematic reviews, TAR researchers are interested in finding a good stabilization point—Finding a point, where the training of the classifier should be stopped. Moreover, in TAR, achieving at least 75% recall is considered as acceptable, whereas in Systematic reviews 95%-100% recall is desirable.

III. Method

For abstract screening, we have title (T_i) and abstract (A_i) for a set of n citations, $\mathcal{C} = \{C_{T_i,A_i}\}_{1\leq i\leq n}$. We represent each citation, C_{T_i,A_i} as a tuple $\langle \mathbf{x}_{C_i}, y_{C_i}\rangle$: $\mathbf{x}_{C_i}$ is a d-dimensional feature space representation of the citation and $y_{C_i} \in \{+1, -1\}$ is a binary label denoting whether C_i is relevant or not. Additionally, we use the following notations throughout the paper. We use $\mathcal{L}$ and $\mathcal{U}$ for the labeled and unlabeled set of citations, respectively; $\mathcal{F}()$ to represent features for a set of citations; h for the hypothesis or hyperplane learned by training on $\mathcal{L}$; and $l(w)$ for the latent space representation for a particular word, w. In this section, we describe the feature space representation and the methods for evaluation of automated abstract screening system.

A. Feature Space Representation

We use two types of feature space representation: (1) Uni-Bigram and (2) Word2Vec. For Word2Vec [48], we train the model on the entire set of citations (abstracts and titles) available in an existing systematic review platform. We use Gensim [49] package with the following parameters: the number of context words as 5, the min word count as 15 and the number of dimensions in the latent space as 500. Thus, for each word in the set of all available words ($w_i \in \mathcal{W}$), we learn a 500-dimensional latent space representation. Some query results from the trained model are shown in the Table II. We can conclude that the model captures the latent similarities among various bio-medical terms. For example, the cosine similarity between "liver" and "cirrhosis" is 0.63, and for the query "The way breast and cancer is related, which words are related to "cirrhosis" in the same way", we can see "liver" is listed as one of the top-5 answers. To represent a citation C_i in the latent space, we use a simple averaging technique to generate a 500 dimensional feature vector, $\mathbf{x}_{C_i}$.

$$\mathbf{x}_{C_i} = \frac{\sum_{w\in\mathcal{W}_{C_i}} l(w)}{|\mathcal{W}_{C_i}|}$$

After averaging, we apply two kinds of normalization: (1) instance normalization (row normalization) and (2) feature based normalization (column normalization). These normalizations give statistically significantly better results than using raw features for threshold agnostic metrics such as AUC and AUPRC. In instance normalization, we z-normalize extracted feature for each citation, $\mathbf{x}_{C_i}$. For column normalization, we z-normalize in each of the 500 dimensions. After both normalizations, we keep the feature values up to 2-decimal places to minimize the memory requirement.

Apache Lucene A standard analyzer [2] is applied to tokenize titles and abstracts and also to generate uni-bigrams. Subsequently, the word counts are binarized. After this step, we obtain a feature space representation for both the labeled and unlabeled citations, i.e. $\mathcal{F}(\mathcal{L})$ and $\mathcal{F}(\mathcal{U})$. Note that the Uni-Bigram is a sparse feature representation whereas Word2Vec

[2]https://lucene.apache.org/core/

Query	Result
model.similarity ('liver', 'cirrhosis')	0.63
model.similarity('breast','cancer')	0.46
model.most_similar(positive=['liver'], top-k with k=10)	hepatic (0.67), cirrhosis (0.63), cholestatic (0.51), spleen (0.51), steatosis (0.5), kidney (0.50), steatohepatitis (0.50), extrahepatic (0.46), pancreas (0.45), cirrhotic (0.45)
model.most_similar(positive=['cancer','cirrhosis'], negative=['breast'], top-k with k=5)	cirrhotic (0.46), hcc (0.46), liver (0.45), heptocellular (0.43), metavir (0.43)

TABLE II
SOME INTERESTING QUERY RESULTS FROM OUR TRAINED WORD2VEC MODEL (THE VALUE INSIDE THE BRACKET REPRESENTS THE COSINE SIMILARITY BETWEEN TWO TERMS.)

Feature Space	Algorithm, Parameter & Loss Function	Method Id
UniBi	SVM^{perf} (B0, AUC)	1
	— (B1, AUC)	2
	— (B1, KLD)	3
	— (B1, QuadMean)	4
	SVM (Default)	5
	SVM^{cost} (J, B0)	6
	— (J, B1)	7
	SVM TRANSDUCTION	11
WORD2VEC ROW	SVM^{perf} (B1, AUC)	21
	— (B1, KLD)	22
	— (B1, QuadMean)	23
	SVM^{cost} (J, B0)	24
	— (J, B1)	25
WORD2VEC COL	SVM^{perf} (B1, AUC)	31
	— (B1, KLD)	32
	— (B1, QuadMean)	33
	SVM^{cost} (J, B0)	34
	— (J, B1)	35

TABLE III
DESCRIPTION AND IDENTIFIER FOR THE SYSTEMS WE HAVE USED FOR EVALUATION PURPOSES. WE USE DEFAULT REGULARIZATION IN ALL OF THE CASES.

is a dense feature representation. Therefore, in general, Uni-Bigram features takes less memory space than the Word2Vec features.

B. Algorithms

In Table III, we outline the different algorithms we used in our comparison along with their parameters, loss functions, and the identifier to represent each algorithm. Table VI uses the identifiers as opposed to the name of the methods. For example, the first row in Table III refers to a method, identified by Id 1, that uses SVM^{perf} with $b = 0$ and AUC as the loss function. We use three types of SVM: (1) Inductive (2) Transductive and (3) SVM for Multivariate Performance. Inductive SVM learns a hypothesis, h induced by $\mathcal{F}(\mathcal{L})$. Transductive SVM answers the most general question: What information do we get from studying $\mathcal{F}(\mathcal{U})$ and how can we use it? It also uses the idea of equivalence classes: two functions from the hypothesis space $\mathcal{H}$ belong to the same equivalence class if they both classify the instances from $\mathcal{F}(\mathcal{U})$ and $\mathcal{F}(\mathcal{L})$ in the same way. This reduces the learning problem of finding a $h \in \mathcal{H}$ to a different learning problem where the goal is to find one equivalence class from infinitely many induced by all the instances in $\mathcal{F}(\mathcal{U})$ and $\mathcal{F}(\mathcal{L})$. On the other hand, SVM^{perf} is an implementation of the Support Vector Machine (SVM) formulation for optimizing multivariate performance measures [19]. It exploits the alternative structural formulation of the SVM optimization problem for conventional binary classification with error rate [50]. We use three different loss functions for the SVM^{perf} implementation:

$$\begin{aligned} \text{AUC} &= \quad \text{\# of swapped pairs} \\ \text{KLD}(p,\hat{p}) &= \quad \sum_{l\in\mathcal{L}} p(l)\log\frac{p(l)}{\hat{p}(l)} \\ \text{QuadMean Error} &= \quad \sqrt{\frac{\frac{\text{FN}}{\text{TP + FN}}\frac{\text{FN}}{\text{TP + FN}} + \frac{\text{FP}}{\text{FP + TN}}\frac{\text{FP}}{\text{FP+TN}}}{2}} \end{aligned} \tag{1}$$

For KLD (*Kullback-Leibler Divergence*) loss function, $p(l)$ and $\hat{p}(l)$ are estimated as follows:

$$\begin{aligned} p(l) &= \quad \frac{\text{TP + FN}}{\text{TP + FP + FN + TN}} \\ \hat{p}(l) &= \quad \frac{\text{TP + FP}}{\text{TP + FP + FN + TN}} \end{aligned}$$

C. Tuning the cost parameter and setting the threshold

For the inductive and Transductive SVM methods, the default cost parameter, denoted as c, is computed as follows:

$$\begin{aligned} b &= \quad \frac{\sum_{1\le i\le|\mathcal{L}|}\|x_{C_i}\|}{|\mathcal{L}|} \\ c &= \quad \frac{1.0}{b*b} \end{aligned}$$

The summation of all the 2-norm of the feature vectors are divided by the number of instances in $\mathcal{F}(\mathcal{L})$ to generate b. Finally, the fraction $\frac{1.0}{b*b}$ gives c. J in SVM^{cost} is set to the following ratio: $\frac{\text{\# of irrelevant}}{\text{\# of relevant}}$. The ratio biases the hypothesis learner (train algorithm) to penalize mistakes on the minority class J (relevant class) times more than the majority class. Furthermore, we follow the recommendation for error loss function from [19] to set the cost parameter for SVM^{perf}:

$$SVM^{perf}(c) = \quad \frac{\text{SVM (c)}\times\text{\# of instances}}{100.0}$$

The motivation of using these values is to make our results more reproducible. To evaluate the reproducibility of each method is one of the main goals of this paper. The way we

calculate the parameter values can also avoid the need for a representative validation set which is very hard to obtain in a systematic review platform. Note, for all the methods, we set the classification threshold to 0.0 (the threshold is very important to separate the relevent citations from the irrelevant ones).

IV. Experimental Results

We run all of our experiments on a computer with a 24-core Intel XEON X5650 2.6Ghz processor running CentOS 6.7 operating system. For each dataset (described in the next section), we perform a 500×2 fold cross validation for each method to analyze the variability. In $n \times k$ fold cross validation, we split the data in k blocks; k-th block becomes the test block and the rest becomes the training data and we repeat this process n times. The split is carried out through stratification. So, for 500×2, we split each dataset into two blocks and then use each block once as a training and once as a test. We repeat this process for 500 times.

A. Datasets

We use 61 reviews for our experiments. Among the 61 reviews, 15 reviews are publicly available from [16] and the rest of the 46 reviews are collected from a deployed platform for producing systematic reviews. In Table III-C, all publicly available reviews start with C (Cohen) whereas reviews from our system starts with P (Private). In the table, we give three statistics about the review—the total number of citations, the total number of relevant citations, and prevalence (ratio of relevant and total number of citations). We divide the reviews into three prevalence groups: (1) Low (0.22% to 5.92%) (2) Mid (6.79% to 13.07%) and (3) High (13.45% to 40.08%). This allows us to study the behavior of the performance metrics, discussed later, within each group. This also helps us present our findings[3] in a detailed fashion.

The datasets of each group are sorted by their prevalence. For example, in dataset P18, the prevalence is 0.22% as 5 (no. of relevant citations) out of 2241 (total citations) is 0.0022.

B. Performance Metrics

To be self contained, we reproduce the metric table from [6]. Table V gives the definition and formula for all the metrics used in the evaluation. The first four measures (Recall, Precision, F-measure and Accuracy) depend on a threshold and are widely used for evaluating automated abstract screening methods. ROC and AUPRC do not depend on thresholds and are common in binary classification problems with data imbalance. We also add two more metrics: (1) Arithmetic mean error and (2) Quadratic mean error. Both of these errors depend on the loss in Recall of the relevant (L_{R_p}) and irrelevant class (L_{R_n}). The arithmetic mean error is the arithmetic mean of L_{R_p} and L_{R_n}. The quadratic mean error is the square root of the arithmetic mean of the squares of L_{R_p} and L_{R_n}. In active learning settings, Burden, Utility, and Yield are frequently used.

[3]For more detailed results:
https://drive.google.com/open?id=0B1RWfbRmk3gVc0ViYW5wX2VSQ28

C. Statistical Test

We have 18 different methods (Table III) and three prevalence groups: (1) Low-prevalence, (2) Mid-prevalence, and (3) High prevalence. We want to compare the methods based on the datasets of a particular prevalence group on a specific metric. Our goal is to generalize the findings on a larger population of possible datasets which fall within a defined prevalence group. So, we model the problem similar to [51]:

$$\text{METRIC} \sim \text{DATA} + \text{METHOD} \tag{2}$$

where DATA has been modeled as random effects and METHOD as a fixed effect. To be more specific, the model resembles $y = mx + c$, where x is a variable and m, c are the constants. In this model, DATA and METHOD mimic x and c respectively. We fit the model in a linear regression framework and perform a two-factor (DATA and METHOD) analysis of variance. This helps us identify whether there is any statistically significant difference among the methods, the datasets and the interactions between the methods and the data. If the test shows statistically significant performance difference among the methods, we again compare them pairwise using paired t-tests (post-hoc testing). Post-hoc testing with 18 methods leads to 18*17 possible post-hoc comparisons. Each of these 18*17 tests needs correction for multiple testing.

To avoid the large number of testing, we follow these steps:

1) We aim to find the best method and compare it with the rest via paired t-tests.
2) All the methods (say, 4) that are not statistically significant (with $\alpha = 0.5$) with respect to the best method falls in the same group. So, we can skip these methods (4) along with the best one for the next comparison.
3) Go to 1. (We repeat the same steps with the remaining 18 (# of methods) - 5 (4 + the best method) =13.)

D. Performance comparison

In Table VI, we present our results (in top three ranked tiers/groups), in terms of the metrics that we define in Table V, for all methods on all datasets. Datasets are grouped by their prevalence, as we explained earlier. The results, for each metric and prevalence, are presented as equivalence groups based on the statistical test (α=0.05). All methods that use the UniBi, WORD2VEC ROW, and WORD2VEC COLUMN features have an identifier as $1-11$ (11 is specifically for SVM Transduction), $21-25$ and $31-35$ respectively (Table III).

We first analyze the metrics which depend on a threshold. For Precision (Prec), cost sensitive SVM (7, where 7 has been used as an identifier for cost sensitive SVM in Table III) outperforms others in the low and mid prevalence groups. But, for the high prevalence group, it falls in the second tier (Table VI). SVM Default (5) also performs well. For Recall, WORD2VEC ROW with SVM^{perf} (AUC) (21) is in the top position in all three prevalence groups. SVM Transduction with Uni-Bi (11) and cost-sensitive SVM with WORD2VEC ROW (25) are the top performers for the $F1$ metric. For ACC, cost-sensitive SVM(B1) with UniBI (7) performs the best followed by SVM Transduction with UniBi (11).

Prevalence $[0.22\% - 5.92\%]$				Prevalence $[6.79\% - 13.07\%]$				Prevalence $[13.45\% - 40.08\%]$			
Review	Total	Pos	Prev (%)	Review	Total	Pos	Prev (%)	Review	Total	Pos	Prev (%)
P18	2241	5	0.22	P39	1149	78	6.79	P41	3250	437	13.45
P31	3034	16	0.53	P16	484	34	7.02	P22	1352	193	14.28
C12	1643	9	0.55	P32	895	63	7.04	P21	1352	193	14.28
P5	8812	60	0.68	P19	541	39	7.21	P36	449	66	14.70
C9	1914	15	0.78	P44	643	50	7.78	P12	820	121	14.76
P30	1864	19	1.02	P23	565	49	8.67	P23	910	137	15.05
P35	2601	33	1.27	P9	257	23	8.95	P20	2703	410	15.17
C1	2544	41	1.61	C6	1113	100	8.98	P17	1704	266	15.61
P7	417	8	1.92	P27	1243	114	9.17	P38	1386	221	15.95
C5	1965	42	2.14	P6	2539	235	9.26	C4	899	146	16.24
C2	845	20	2.37	P28	1242	115	9.26	P1	906	150	16.56
C13	3377	85	2.52	P43	996	95	9.54	P24	2488	419	16.84
P11	1580	41	2.59	P15	954	95	9.96	P42	4019	715	17.79
C14	660	24	3.64	P10	616	63	10.23	P2	1484	265	17.86
P26	351	13	3.70	P33	640	68	10.63	P8	498	100	20.08
C11	1330	51	3.83	P46	551	59	10.71	C7	368	80	21.74
P4	1187	56	4.72	P34	1728	200	11.57	P45	957	230	24.03
C3	296	16	5.41	C8	343	41	11.95	C10	503	136	27.03
P37	730	40	5.48	P40	850	110	12.94	P3	1487	404	27.17
P25	338	20	5.92	C15	306	40	13.07	P14	822	256	31.14
								P13	819	328	40.08

TABLE IV
DATASET STATISTICS. ALL PUBLICLY AVAILABLE REVIEWS START WITH C (COHEN) WHEREAS REVIEWS FROM OUR SYSTEM STARTS WITH P (PRIVATE)

Now, we analyze the performance of the methods on two threshold agnostic metrics. In AUC, WORD2VEC ROW with SVM^{perf} (AUC) (21) is the one with the best performance. Cost sensitive SVM(B1) with WORD2VEC ROW normalized feature (25) also performs well. In AUCPR, SVM^{perf} with (B0, AUC) and (B1, AUC) along with the Uni-BI feature are the top performing methods.

Cost sensitive SVM with WORD2VEC ROW normalized feature, 24 and 25, rank first in the AM and QUADMEAN ERRORS metrics, respectively. Again WORD2VEC ROW with SVM^{perf} (AUC) (21) ranks first in the YIELD and UTILITY metrics while the SVM Default (5) tops the list in the BURDEN metric.

In conclusion, we observe that that there is almost always a method that ranks first in the three prevalence groups. However, there is no single dominant method across all metrics. In summary, there is no single "winner" or best method—various methods perform well on different prevalence groups and for different metrics. For instance, WORD2VEC ROW with SVM^{perf} (AUC) (21) seems to be a good choice, outperforming the other methods in five metrics. However it is not present in any of the equivalence groups for a few other metrics. This is because the method has a very high recall, but very low precision. The low precision might come from the averaging technique (Section III-A) we use to represent the feature space. It follows that a holistic "composite" strategy must be adopted in any practical system. We show such a method in Section V. This is not surprising since each method has been designed to meet a certain objective. However, abstract screening is a complicated process, in which, a "good" method should be able to optimize several metrics such as precision, recall, AUC, utility and possibly many others simultaneously.

E. Experiment in an active learning setting

In a practical system, users will usually label citations by batches. Hence, an active learning approach has the potential to achieve better results [12]. We use the SVM^{perf} (AUC) (21) method for this experiment as it is the top performing method in terms of the AUC metric in all three prevalence groups (Table VI). Initially for training, we choose 5 relevant and 45 irrelevant citations uniformly at random from the entire set and then learn a hyperplane h. We calculate the distance (score) from the hyperplane for each of the unlabeled instances and rank them based on this score. We choose the top-50 from the ranked citations to retrain the model along with the existing labeled citations. We repeat this experiment 500 times and take the average. The goal of this experiment is to see the following: if a particular user labels 50 top ranked citations per batch, what is the percentage of total citations that the user has to screen to get all the relevant ones?

Figure 1 shows the results of the proposed experiment. The results are based on three groups of prevalence values. For the low prevalence group, out of 20 reviews, 7 reviews (the first four values on the top histogram) need 40% of the total citations to be screened to get all the relevant citations, 12 need around 50 to 70% citations, and 1 needs more than 90%. For the mid and high prevalence groups, 9 out of 20 and 11 out of 21 reviews need around 80% to 90% citations to be screened to get all the relevant citations. In the bottom row of Figure 1, we show the spread of standard deviations for each of the prevalence groups. As expected, the low prevalence

Metric	Definition	Formula
Recall (Sensitivity)	Ratio of correctly predicted relevant citations to all relevant ones.	$\frac{\text{TP}}{\text{TP + FN}}$
Precision	Ratio of correctly identified relevant citations to all of those predicted as relevant.	$\frac{\text{TP}}{\text{TP + FP}}$
F-Measure	Combines Precision and Recall values. It corresponds to the harmonic mean of Precision and Recall for $\beta = 1$.	$\frac{(1+\beta^2).\text{Precision.Recall}}{\beta^2.\text{Precision + Recall}}$
Accuracy	Ratio of relevant and irrelevant citations predicted correctly to all citations.	$\frac{\text{TP + TN}}{\text{TP + TN + FP + FN}}$
ROC (AUC)	Area under the curve obtained by graphing the true positive rate against the false positive rate; 1.0 is a perfect score and 0.5 is equivalent to a random ordering.	
AUPRC	Area under precision recall curve.	
AM ERROR	Arithmetic mean of the loss in Recall of the relevant (L_{R_p}) and irrelvant class (L_{R_n})	$\frac{\frac{\text{FN}}{\text{TP + FN}} + \frac{\text{FP}}{\text{FP + TN}}}{2}$
QUADMEAN ERROR	Quadratic mean, aka. root mean square, measures the magnitude of varying quantities. It is defined as the square root of the arithmetic mean of the squares of the loss in Recall of the relevant (L_{R_p}) and the irrelevant class (L_{R_n}).	$\sqrt{\frac{\frac{\text{FN}}{\text{TP + FN}} \cdot \frac{\text{FN}}{\text{TP + FN}} + \frac{\text{FP}}{\text{FP + TN}} \cdot \frac{\text{FP}}{\text{FP + TN}}}{2}}$
Burden	The fraction of the total number of citations that a human must screen.	$\frac{\text{TP}^L + \text{TN}^L + \text{TP}^U + \text{FP}^U}{\text{N}}$
Yield	The fraction of citations that are identified by a given screening approach.	$\frac{\text{TP}^L + \text{TP}^U}{\text{TP}^L + \text{TP}^U + \text{FN}^U}$
Utility	Utility is a weighted sum of Yield and Burden. Here, β is a constant. It represents the relative importance of Yield, in comparison to Burden. We use $\beta = 19$ in our experimental evaluations following the suggestion from [18].	$\frac{\beta.\text{Yield} + (1 - \text{Burden})}{\beta + 1}$

TABLE V
VARIOUS METRICS WITH THEIR FORMULA AND DEFINITIONS. TP, FP, TN AND FN REPRESENTS TRUE POSITIVE, FALSE POSITIVE, TRUE NEGATIVE AND FALSE NEGATIVE RESPECTIVELY.

group has a very high standard deviation. For the mid and high prevalence groups, the standard deviation is very small (around $\sim$ 0.01-0.05) in most of the cases.

In Figure 2, we further analyze the inclusion behavior of a particular random run of review 1. As the results are similar for some of the other reviews, we omit the plots to avoid repetitions. For a random run, around 58%, 1500 out of 2544 documents have to be screened before it gets its final relevant citation. However, we clearly see from the figure that, it gets almost all but the final one after screening only 400 out of 2544 which is around 15% of the total citations. The last citations surprisingly takes 22 more iterations. With the help of a domain expert, it would be interesting to see whether the final citation has any outlying characteristics.

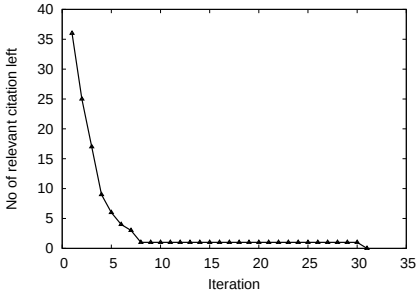

Fig. 2. The pattern of decrease in the number of relevant citations during abstract screening for a particular review.

Algorithm 1: RelRank: A Five Star rating algorithm using ensemble of max-margin based methods

Input : $\mathcal{L}$, Labeled dataset; $\mathcal{U}$, Unlabeled dataset
Output: Score, $\mathcal{S}_{1 \leq i \leq |\mathcal{U}|}$

$\mathcal{F}_{\mathcal{L}}$, $\mathcal{F}_{\mathcal{U}} \leftarrow$ **GenerateFeature** ($\mathcal{L}$, $\mathcal{U}$, feature = W)
$h1 \leftarrow$ **Train** (SVM^{perf}, $\mathcal{F}_{\mathcal{L}}$)
$h2 \leftarrow$ **Train** (SVM^{cost}, $\mathcal{F}_{\mathcal{L}}$)
$\mathcal{S}_{h1} \leftarrow$ **Predict** ($h1$, $\mathcal{F}_{\mathcal{U}}$)
$\mathcal{S}_{h2} \leftarrow$ **Predict** ($h2$, $\mathcal{F}_{\mathcal{U}}$)
$\mathcal{F}_{\mathcal{L}}$, $\mathcal{F}_{\mathcal{U}} \leftarrow$ **GenerateFeature** ($\mathcal{L}$, $\mathcal{U}$, feature = U)
$h3 \leftarrow$ **Train** (SVM^{cost}, $\mathcal{F}_{\mathcal{L}}$)
$\mathcal{S}_{h3} \leftarrow$ **Predict** ($h3$, $\mathcal{F}_{\mathcal{U}}$)
$\mathcal{S} \leftarrow$ **GenerateCombinedScore** ($\mathcal{U}$, $\mathcal{S}_{h1}$, $\mathcal{S}_{h2}$, $\mathcal{S}_{h3}$)
return $\mathcal{S}$

V. PROPOSED 5-STAR RATING ALGORITHM

In our systematic review platform, we wanted to rank citations based on their graded relevance. The intuition behind this proposal is to help reviewers better manage their time. For this purpose, we rate the citations from $1-5$ using **RelRank** (described later in Algorithm 1). The citations with 5 star might be relevant (with high probability) and need more attention whereas 1 starred documents might be irrelevant (with high probability) and may need less attention. Among the 5 stars, we conceptualize $3-5$ star to indicate relevant citations and $1-2$ star for irrelevant ones.

	Prevalence $[0.22\% - 5.92\%]$			Prevalence $[6.79\% - 13.07\%]$			Prevalence $[13.45\% - 40.08\%]$		
Metric	$rg = 1$	$rg = 2$	$rg = 3$	$rg = 1$	$rg = 2$	$rg = 3$	$rg = 1$	$rg = 2$	$rg = 3$
PREC	**7**, 11	3, **5**, 25	4, 6, 35	**5**,**7**	11	3, 6, 22, 25	**5**	**7**	11
REC	**21**	31	1, 2, 34	**21**	1, 2, 31, 34	22	**21**	1,2	31
F1	**11**, 25	24, 35	4, 22, 23	**11**, 25	4, 22, 23, 24, 35	3, 6, 33	7, **11**, 25	35	24
ACC	**7**	5	3, 6, 11	5, **7**	11	6	5, **7**	11	3, 4, 6
AUC	**21**, **25**	31, 35	1, 2, 5, 7, 23, 24	**21**, **25**	4, 24, 31, 35	1, 2, 5, 7, 22, 23	**21**	7, 25	(All Others)
AUPRC	**1**, **2**, 5, 7, 21	4, 25	11, 23, 24, 31	**1**, **2**, 4, 5	7, 21	25	**1**, **2**	(All Others)	-
AM ERROR	**4**, 31, 35, 24, **25**	1, 2, 11, 22, 23, 32, 33, 34	3, 6, 7, 21	**4**, **25**	24	31, 33, 35	**25**	24	4, 35
QUADMEAN ERROR	1, 2, 4, **24**, 25, 31, 32, 33	23, 34, 35	11, 22 4, 25, 33	**24**	23, 31, 35	25	**24**	4, 35	
YIELD	**21**	31	1, 2, 34	**21**	1, 2, 31, 34	32	**21**	1, 2	31
BURDEN	**5**	7	3, 6	**5**	7	6	**5**, 7	11	6
UTILITY	**21**	31	1, 2, 34	**21**	1, 2, 31, 34	32	**21**	1,2	31
AUC (STD)	**21**, 22, 23, 25, 31, 35	All Others		**21**, 25	24, 31, 35	22, 23	**21**	25, 31	35
AUPRC (STD)	**32**, 34	3, 6, 22, 33	(All others)	No statistically significant difference among methods			No statistically significant difference among methods		

TABLE VI
GROUPING RESULTS OF ALL THE SYSTEMS BASED ON VARIOUS METRICS IN TABLE V (rg STANDS FOR RANK GROUP).

We present our 5-star rating algorithm, **RelRank** using an ensemble of SVM based methods. We choose method 21, 25, and 7 (See Table III) because of their special characteristics. For instance, Method 21 outperforms others based on the AUC and RECALL metric. Moreover, the Method 21 has the lowest standard deviation for the AUC metric (See Table VI). While evaluated on the F1 metric which is a harmonic mean of Precision and Recall, Method 25 produces the highest F1 measure. Method 11 has similar effect as Method 25 but with much higher computation time. On the other hand, Method 7 has the highest Precision. So, in **RelRank** , if all of these three methods (Methods 21, 25 and 7) agree on the relevance of a citation, then the citation gets a 5 star; if two of them agrees then it gets a 4 star. Within a particular star, citations are ranked based on their average ranks (described in Algorithm 2).

In **RelRank** , we first generate WORD2VEC ROW features for both the $\mathcal{L}$ and $\mathcal{U}$. Then, we train SVM^{perf} and SVM^{cost} (Methods 21 and 25) to generate first ($h1$) and second hyperplane ($h2$) respectively. Step 4 and 5 generate scores for $\mathcal{U}$ based on the distances from $h1$ and $h2$. Using the Uni-Bi gram features (Step 6), we compute a third score $\mathcal{S}_{h3}$ using hyperplane $h3$ (Step 7). In Step 9, we combine the three scores to generate a final score for $\mathcal{U}$ (which we formally present as Algorithm 2. Please see the appendix for more details).

We evaluate our 5-star algorithm on each of the citations in different cumulative star-groups. The results are provided in Table VII. For instance, we evaluate the citations which get 3 or more stars, 4 or more stars and 5 stars. We observe that our 5-star algorithm performs similarly to Method 21 for the citations receiving 3 stars or above. However, for the precision metric, the 5-star algorithm performs better than Method 21 as it falls in the rank group 7 whereas Method 21 falls in the group 9. This is because the 5-star algorithm is a combination of Method 21 with two other methods that have better precision. However, for the recall metric, **RelRank** (3-star) are in the top group whereas **RelRank** (4-star) and **RelRank** (5-star) are in the second top group. This is obvious as increasing recall might lose precision. We observe a similar behavior in the remaining metrics.

VI. CONCLUSIONS AND FUTURE WORK

Automating the production of systematic reviews will be crucial in delivering the promise of evidence-based medicine. An key task in this process is sifting through hundreds to thousands of citations and identifying relevant studies for further analysis. In this paper, we studied the most popular classification methods employed in this task. We focus on the methods that better fit the constraints of a practical system for abstract screening. In all, we report on 18 methods on a

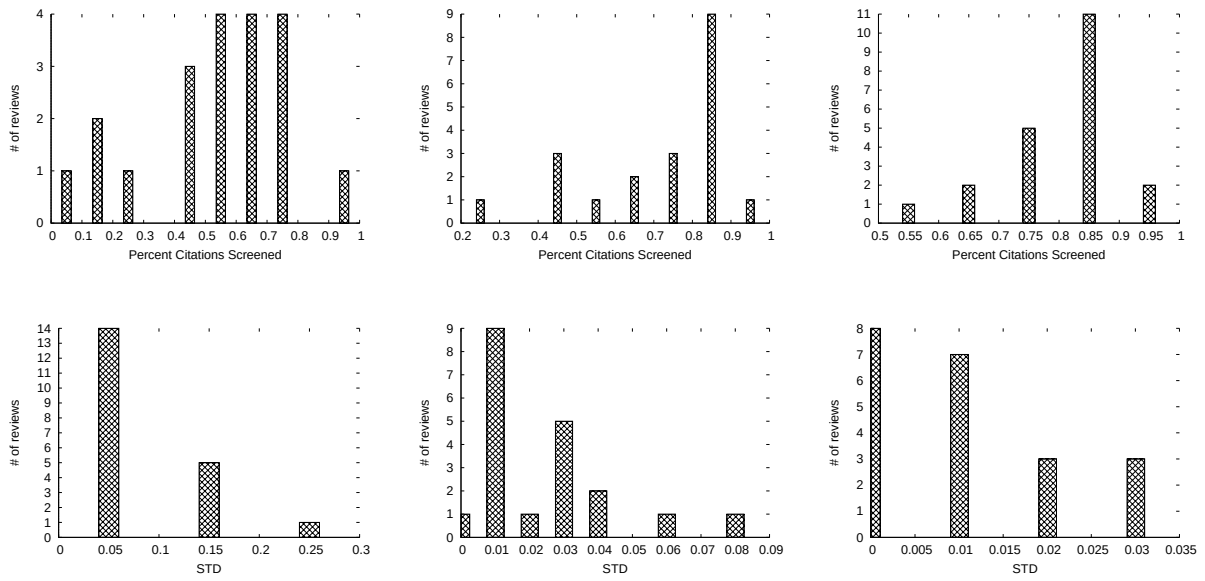

Fig. 1. Active Learning experiment with variable test set

Metric	RelRank (3-star)	RelRank (4-star)	RelRank (5-star)
Recall (Sensitivity)	$\langle \mathbf{1}, 1, 1 \rangle$	$\langle \mathbf{2}, 2, 2 \rangle$	$\langle \mathbf{3}, 4, 5 \rangle$
Precision	$\langle 9, 10, 12 \rangle$	$\langle 8, 9, 10 \rangle$	$\langle \mathbf{7}, 7, 7 \rangle$
F-Measure	$\langle 7, 7, 9 \rangle$	$\langle 6, 6, 7 \rangle$	$\langle 5, 4, 4 \rangle$
Accuracy	$\langle 9, 13, 11 \rangle$	$\langle 8, 11, 9 \rangle$	$\langle 7, 9, 6 \rangle$
ROC (AUC)	$\langle 1, 1, 1 \rangle$	$\langle 1, 1, 1 \rangle$	$\langle 1, 1, 1 \rangle$
AUPRC	$\langle 1, 1, 1 \rangle$	$\langle 1, 1, 1 \rangle$	$\langle 1, 1, 1 \rangle$
AM ERROR	$\langle 5, 7, 8 \rangle$	$\langle 3, 5, 4 \rangle$	$\langle 1, 1, 1 \rangle$
QUADMEAN ERROR	$\langle 5, 9, 9 \rangle$	$\langle 3, 6, 6 \rangle$	$\langle 1, 1, 1 \rangle$
Burden	$\langle 11, 14, 13 \rangle$	$\langle 10, 13, 11 \rangle$	$\langle 8, 10, 8 \rangle$
Utility	$\langle 1, 1, 1 \rangle$	$\langle 2, 2, 2 \rangle$	$\langle 3, 4, 5 \rangle$

TABLE VII
PERFORMANCE OF RELRANK IN DIFFERENT STARRED THRESHOLD. LESSER THE RANK GROUP BETTER THE PERFORMANCE IN A PARTICULAR METRIC. THREE NUMBERS INSIDE THE BRACKET INDICATES RANK GROUP FOR ⟨ LOW, MID, HIGH ⟩ PREVALENCE GROUP.

very large number of reviews (61) using 11 metrics. There is no single "winner" or best method—various methods perform well on different prevalence groups and for different metrics. For instance, WORD2VEC ROW with SVM^{perf} (AUC) seems to be a good choice, outperforming the other methods in five metrics but is not doing well for a few other metrics. We also observe that in an active learning setting, a very large portion of included citations can be easily found with few iteration. However, one or two citations have some outlying behavior and requires much more iterations to be found.

We also presented an ensemble method that combines three of the methods we evaluated and convert their scores into a 5-star rating system through a voting mechanism and rank citations based on their graded relevance. The goal is to help reviewers better manage their time; A 5-star citation might be relevant with high probability and need more attention whereas a 1-star citation might be irrelevant with high probability and thus need less attention.

As a future work, we plan to look at different approaches to combine methods as well look more carefully at the outlying citations and how they can be screened with less iterations.

REFERENCES

[1] T. C. Chalmers, H. Smith, B. Blackburn, B. Silverman, B. Schroeder, D. Reitman, and A. Ambroz, "A method for assessing the quality of a randomized control trial," *Controlled clinical trials*, vol. 2, no. 1, pp. 31–49, 1981.

[2] Sense-About-Science. (2009, November) Sense about systematic reviews. [Online].

[3] L. S. Uman, "Systematic reviews and meta-analyses," *Journal of the Canadian Academy of Child and Adolescent Psychiatry*, vol. 20, no. 1, p. 57, 2011.

[4] S. Agarwal and D. Roth, "Learnability of bipartite ranking functions," in *Learning Theory*. Springer, 2005, pp. 16–31.

[5] M.-F. Balcan, N. Bansal, A. Beygelzimer, D. Coppersmith, J. Langford, and G. B. Sorkin, "Robust reductions from ranking to classification," *Machine learning*, vol. 72, no. 1-2, pp. 139–153, 2008.

[6] A. OMara-Eves, J. Thomas, J. McNaught, M. Miwa, and S. Ananiadou, "Using text mining for study identification in systematic reviews: a systematic review of current approaches," *Systematic reviews*, vol. 4, no. 1, p. 1, 2015.

[7] T. Raeder, *Evaluating and maintaining classification algorithms*. University of Notre Dame, 2012.

[8] T. Raeder, T. R. Hoens, and N. V. Chawla, "Consequences of variability in classifier performance estimates," in *2010 IEEE 10th International Conference on Data Mining (ICDM)*, 2010, pp. 421–430.

[9] K. Ambert, "A prospective evaluation of an automated classification system to support evidence-based medicine and systematic review," vol. 2010, p. 121, 2010.

[10] Y. Mo, G. Kontonatsios, and S. Ananiadou, "Supporting systematic reviews using lda-based document representations," *Systematic reviews*, vol. 4, no. 1, p. 1, 2015.

[11] M. Khabsa, A. Elmagarmid, I. Ilyas, H. Hammady, and M. Ouzzani, "Learning to identify relevant studies for systematic reviews using random forest and external information," *Machine Learning*, pp. 1–18, 2015.
[12] M. Miwa, J. Thomas, A. OMara-Eves, and S. Ananiadou, "Reducing systematic review workload through certainty-based screening," *Journal of biomedical informatics*, vol. 51, pp. 242–253, 2014.
[13] A. M. Cohen, K. Ambert, and M. McDonagh, "Studying the potential impact of automated document classification on scheduling a systematic review update," *BMC medical informatics and decision making*, vol. 12, no. 1, p. 1, 2012.
[14] B. K. Olorisade, E. De Quincey, O. Brereton, P. Andras *et al.*, "A critical analysis of studies that address the use of text mining for citation screening in systematic reviews," 2016.
[15] A. M. Cohen, "Optimizing feature representation for automated systematic review work prioritization," in *AMIA annual symposium proceedings*, vol. 2008, 2008, p. 121.
[16] A. M. Cohen, W. R. Hersh, K. Peterson, and P.-Y. Yen, "Reducing workload in systematic review preparation using automated citation classification," *Journal of the American Medical Informatics Association*, vol. 13, no. 2, pp. 206–219, 2006.
[17] S. Hempel, K. D. Shetty, P. G. Shekelle, L. V. Rubenstein, M. S. Danz, B. Johnsen, and S. R. Dalal, "Machine learning methods in systematic reviews: identifying quality improvement intervention evaluations," 2012.
[18] B. C. Wallace, K. Small, C. E. Brodley, and T. A. Trikalinos, "Active learning for biomedical citation screening," in *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2010, pp. 173–182.
[19] T. Joachims, "A support vector method for multivariate performance measures," in *Proceedings of the 22nd international conference on Machine learning*, 2005, pp. 377–384.
[20] ——, "Transductive inference for text classification using support vector machines," in *ICML*, vol. 99, 1999, pp. 200–209.
[21] S. Shalev-Shwartz, Y. Singer, N. Srebro, and A. Cotter, "PEGASOS: Primal estimated sub-gradient solver for svm," *Mathematical programming*, vol. 127, no. 1, pp. 3–30, 2011.
[22] C. C. Aggarwal, *Data classification: algorithms and applications*. CRC Press, 2014.
[23] Y. Sun, H. Deng, and J. Han, "Probabilistic models for text mining," in *Mining Text Data*. Springer, 2012, pp. 259–295.
[24] A. Esuli and F. Sebastiani, "Optimizing text quantifiers for multivariate loss functions," *ACM Transactions on Knowledge Discovery from Data (TKDD)*, vol. 9, no. 4, p. 27, 2015.
[25] W. Liu and S. Chawla, "A quadratic mean based supervised learning model for managing data skewness." in *SDM*, 2011, pp. 188–198.
[26] C. Elkan, "The foundations of cost-sensitive learning," in *International joint conference on artificial intelligence*, vol. 17, no. 1, 2001, pp. 973–978.
[27] Y. Yue, T. Finley, F. Radlinski, and T. Joachims, "A support vector method for optimizing average precision," in *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*, 2007, pp. 271–278.
[28] C. Chen, A. Liaw, and L. Breiman, "Using random forest to learn imbalanced data," *University of California, Berkeley*, 2004.
[29] H. Han, W.-Y. Wang, and B.-H. Mao, "Borderline-SMOTE: a new over-sampling method in imbalanced data sets learning," in *Advances in intelligent computing*, 2005, pp. 878–887.
[30] C. Bunkhumpornpat, K. Sinapiromsaran, and C. Lursinsap, "Safe-level-SMOTE: Safe-level-synthetic minority over-sampling technique for handling the class imbalanced problem," in *Advances in knowledge discovery and data mining*, 2009, pp. 475–482.
[31] B. C. Wallace and I. J. Dahabreh, "Improving class probability estimates for imbalanced data," *Knowledge and Information Systems*, vol. 41, no. 1, pp. 33–52, 2014.
[32] ——, "Class probability estimates are unreliable for imbalanced data (and how to fix them)," in *12th International Conference on Data Mining*, 2012, pp. 695–704.
[33] A. Niculescu-Mizil and R. Caruana, "Predicting good probabilities with supervised learning," in *Proceedings of the 22nd international conference on Machine learning*, 2005, pp. 625–632.
[34] S. Tong and D. Koller, "Support vector machine active learning with application to text classification," vol. 2, 2001, pp. 45–66.
[35] S. Hanneke, "Theory of active learning," 2014.
[36] I. Dagan and S. P. Engelson, "Committee-based sampling for training probabilistic classifiers," in *Proceedings of the Twelfth International Conference on Machine Learning*, 1995, pp. 150–157.
[37] T. Scheffer, C. Decomain, and S. Wrobel, "Active hidden markov models for information extraction," in *Advances in Intelligent Data Analysis*. Springer, 2001, pp. 309–318.
[38] A. Culotta and A. McCallum, "Reducing labeling effort for structured prediction tasks," in *AAAI*, 2005, pp. 746–751.
[39] H. S. Seung, M. Opper, and H. Sompolinsky, "Query by committee," in *Proceedings of the fifth annual workshop on Computational learning theory*. ACM, 1992, pp. 287–294.
[40] R. Nowak, "Noisy generalized binary search," in *Advances in neural information processing systems*, 2009, pp. 1366–1374.
[41] J. Fu and S. Lee, "Certainty-enhanced active learning for improving imbalanced data classification," in *Data Mining Workshops (ICDMW), 2011 IEEE 11th International Conference on*. IEEE, 2011, pp. 405–412.
[42] H. L. Roitblat, A. Kershaw, and P. Oot, "Document categorization in legal electronic discovery: computer classification vs. manual review," *Journal of the American Society for Information Science and Technology*, vol. 61, no. 1, pp. 70–80, 2010.
[43] M. Gabriel, C. Paskach, and D. Sharpe, "The challenge and promise of predictive coding for privilege," in *ICAIL 2013 DESI V Workshop*, 2013.
[44] D. W. Henry, "Predictive coding: Explanation and analysis of judicial impact and acceptance compared to established e-commerce methodology," http://www.dwhenry.com/files/Predictive%20Coding.pdf, [Online;Accessed 23-June-2015].
[45] M. R. Grossman and G. V. Cormack, "Technology-assisted review in e-discovery can be more effective and more efficient than exhaustive manual review," *Rich. JL & Tech.*, vol. 17, p. 1, 2010.
[46] C. GLOSSARY, "The grossman-cormack glossary of technology-assisted review," *Federal Courts Law Review*, vol. 7, no. 1, 2013.
[47] T. K. Saha, M. Al Hasan, C. Burgess, M. A. Habib, and J. Johnson, "Batch-mode active learning for technology-assisted review," in *Big Data (Big Data), 2015 IEEE International Conference on*. IEEE, 2015, pp. 1134–1143.
[48] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," in *Advances in neural information processing systems*, 2013, pp. 3111–3119.
[49] R. Řehůřek and P. Sojka, "Software Framework for Topic Modelling with Large Corpora," in *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*. ELRA, 2010, pp. 45–50.
[50] T. Joachims, "Training linear svms in linear time," in *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2006, pp. 217–226.
[51] M. S. Cohen AM, "RMEQ: A tool for computing equivalence groups in repeated measures studies. in: Linking literature," in *Information and Knowledge for Biology: Proceedings of the BioLINK2008 Workshop*, 2008.

APPENDIX

Algorithm 2 helps the already classified citations to preseent in a ranked manner based on the scores (obtained by calculating distance from hyperplanes). Algorithm 2 takes separation threshold (ST) and max range (MR) as a parameter. ST is used to define separation between ratings. We use 1000 as ST. MT is used for range normalization of a rank score between $1 - \text{MR}$. Step 1 gets the number of unlabeled citations. Step 2 and Step 3 define FVOTE and NORM as a lambda function. Both of them are used to transform ranks into a specified range. FVOTE smooths any value between $0.36 - 1.0$.

In Step 4, we get ranks for each of the scores $\mathcal{S}_{1 \leq h \leq |h|}$. The higher the distance is from the hyperplane for a particular citation, the higher the rank. Steps 5-16 are used for calculating the final combined score. We iterate through all the scores in order. The very first hypothesis has the supreme power. If it predicts a particular citation as relevant, then the citation

is rated between 3 to 5 star, and if it predicts as irrelevant, then the same gets a 1 or 2 star. In Step 6, we get the vote from the dominant classifier i.e. if the score is greater than the classifier threshold then the vote count increases by 1. The rank score and the fractional vote are obtained in Steps 7 and 8. The number of vote (nv) only increases if both the dominant classifier and the current hypothesis vote as relevant. On the other hand, the number of vote decreases if the both of them agree on its irrelevancy. Finally, in Step 14, we get a normalize rank between 1 to MR. If the number of votes is greater than or equal to 1 then fractional vote is considered as the final vote otherwise the negative votes are used unchanged. This ensures that even if a particular citation gets postive votes, it still needs to get top rank to maintain its position in the rating. We use 1000, 800, 0.0, 2000 and 2500 as a value for the parameters: ST, MR, threshold for **RelRank** (3-star), **RelRank** (4-star), and **RelRank** (5-star).

Example: For example, if we decide that a 5 starred citation needs to reside in the top $1/3$ of the whole set then it has to get at least 0.66 fractional vote from each hypothesis. It gets a score of $0.66 + 0.66 + 0.66 = 19.8$. If the separation threshold is 1000 then it recieves a score of $1.98 * 1000 = 1980$. From the normalized score, it gets $2/3 * 800(MR)$ = 533. So, the final score becomes greater than 2500. Hence, we can set a threshold of 2500 for the 5 star rating. To be more conservative we can set a high threshold. The maximum score a citation can get is $3.0 * 1000 + 800 = 3800$.

Algorithm 2: GenerateCombinedScore

Input : Unlabeled dataset($\mathcal{U}$), $\mathcal{S}_h[]$, $\mathcal{T}_h[]$, Separation Threshold (ST), MAX Range (MR)
Output: Score, $\mathcal{S}$

1 U = $|\mathcal{U}|$
2 FVOTE = **LambaF** $\{|r|\exp(-\frac{\mathrm{U}-r}{\mathrm{U}})\}$
3 NORM = **LambaF** $\{|s|\frac{s-1.0}{\mathrm{U}-1.0} * \mathrm{MR}\}$
4 r_h = **GetRanksForASetOfElements** $(\mathcal{S}_h)$
5 **for** $u = 0..|U|$ **do**
6 nv = $\mathcal{S}_{h[0][u]} \geqslant \mathcal{T}_{h[0]}$? 1 : -1
7 rs = $r_h[0][u]$
8 fv = FVOTE (rs)
9 **for** $v = 1..|h|$ **do**
10 rs_v = $r_h[v][u]$
11 fv = fv + FVOTE (rs_v)
12 rs = rs + rs_v
13 nv = $\mathcal{S}_{h[v][u]} \geqslant \mathcal{T}_{h[v]}?(nv > 0?nv + 1 : nv) : (nv < 0?nv - 1 : nv)$
14 ns = NORM ($rs/|h|$)
15 finv = nv $\geqslant$ 1 ? fv : nv
16 $\mathcal{S}[u]$ = finv * ST + ns
17 **return** $\mathcal{S}$